arXiv:1403.0395v1 [math-ph] 3 Mar 2014

Inverse Problems and Imaging
Volume X, No. 0X, 20xx, X–XX
doi:10.3934/ipi.xx.xx.xx

POINCARÉ INVERSE PROBLEM AND TORUS CONSTRUCTION IN PHASE SPACE

Teemu Laakso and Mikko Kaasalainen

Department of Mathematics, Tampere University of Technology
PO Box 553
33101 Tampere, Finland

(Communicated by the associate editor name)

Abstract. The phase space of an integrable Hamiltonian system is foliated by invariant tori. For an arbitrary Hamiltonian H such a foliation may not exist, but we can artificially construct one through a parameterised family of surfaces, with the intention of finding, in some sense, the closest integrable approximation to H. This is the Poincaré inverse problem (PIP). In this paper, we review the available methods of solving the PIP and present a new iterative approach which works well for the often problematic thin orbits.

1. **Introduction.** Integrability of a dynamical system, witnessed by a foliation of invariant toroidal manifolds in the phase space, is a highly useful property. For a Hamiltonian H_0 which is autonomous and integrable, we have a symplectic map (often explicit through the Hamilton-Jacobi approach)

$$(\theta, J) \leftrightarrow (q, p), \tag{1}$$

between the action-angle variables (θ, J) and the Cartesian phase-space coordinates (q, p). Not only does this eliminate the need for any numerical integration of the equations of motion, but also allows one to study the canon of near-integrable Hamiltonians H in a form of power series;

$$H(\theta, J) = H_0(J) + \varepsilon H_1(\theta, J) + \varepsilon^2 H_2(\theta, J) + \ldots, \tag{2}$$

where the strength of the perturbation ε is small. To study Eq. (2) and the corresponding equations of motion, according to Poincaré [18], is "the general problem of dynamics". We consider it as the direct problem, and ask the inverse: "given a near-integrable Hamiltonian in the form $H(q, p)$, what is the best way of defining the maps (1) and $J \mapsto H_0(J)$?", i.e., we are looking for the closest integrable approximation to H. Indeed, there are many cases where a system shows near-integrable behaviour numerically, but there is no prior information on H_0, underlining the relevance of the question above.

The Poincaré inverse problem (PIP) is obviously an ill-posed problem (in the sense of Hadamard); particularly, in regions of phase space, where invariant toroidal manifolds are absent, the closeness of H and H_0 is an insufficient requirement for a unique solution. Simply minimising $\|H - H_0\|$ everywhere can lead to unacceptable

2010 *Mathematics Subject Classification.* Primary: 70H07, 70H08, 70K43; Secondary: 65P10, 85A05.

Key words and phrases. Near integrability, invariant torus, torus construction, surface construction in N dimensions, Poincare inverse problem, geometric inverse problems.

 ©20xx American Institute of Mathematical Sciences

results [7]. Generally, in order to guarantee that the solution is a reasonable approximation to the actual motion, additional constraints (regularisation) are necessary. On the other hand, since H is a known function, the data are error-free sampling points, and their number is only limited by the available computational resources. Therefore, "model noise" is the fundamental source of error, which is typical for many inverse problems.

The PIP can also be a subproblem of a larger inverse one. For example, in dynamical or phase-space tomography [9] we want to solve for a potential field Φ and phase-space distribution functions best corresponding to observed positions and velocities of matter. In this process, we repeatedly construct an H_0 corresponding to a trial Φ in order to describe the phase-space distribution of observed matter along the tori of H_0.

Our approach towards the solution of the PIP is based on numerical construction of phase-space tori which should coincide with the existing Kolmogorov-Arnold-Moser (KAM) tori of H (if any); i.e., on these, $\|H - H_0\| = 0$ within the available accuracy. This ensures that the tori of $H = \mathcal{H}(\mathcal{J})$, where $\mathcal{H}$ is any specific integrable Hamiltonian and $\mathcal{J}$ are its actions, are faithfully reconstructed. Also, for a perturbed $H = \mathcal{H} + \varepsilon\mathcal{H}_1 + \varepsilon^2\mathcal{H}_2 + \ldots$ of the form (2), we have the natural limit

$$\lim_{\varepsilon\to 0} H_0(J;\varepsilon) = \lim_{\varepsilon\to 0} H(q,p;\varepsilon) = \mathcal{H}(\mathcal{J}).$$

Note that it is not necessary to have $H_0 = \mathcal{H}$, when $\varepsilon > 0$. Even here there is some ambiguity; a torus can be analytically defined as a KAM one only by analysing it via the perturbation (2); whether we can directly deem an orbit of H to correspond to a KAM torus is usually up to numerical resolution. As implied by the above principles, our approach is geometric rather than dynamical (which was the line of thought preferred by Poincaré as well).

We construct the tori one by one, and shall refer to this process as torus construction, torus modelling, or torus embedding. The full solution, defining the maps (1) and $J \mapsto H_0(J)$ everywhere, requires an interpolation scheme between the constructed tori. What interpolations produce a foliated global set of tori that really correspond to an H_0 is an open question. We shall not attempt to answer it in this paper; instead, we concentrate on the fundamental geometric problem of defining the local tori. This part of the PIP reduces to the problem of surface reconstruction, if only the KAM tori are modelled. However, this is generally not sufficient; when H is far from integrable, its set of KAM tori may be too sparse to form an interpolation grid. One may even have a case where H has no KAM tori at all [8]. Hence, the ability of embedding approximating tori in the absence of KAM ones is also of practical importance, and it shall be a priority in our approach.

This paper is organised as follows. In Sect. 2 we review and classify some of the available methods of torus construction, and in Sect. 3 we introduce a new one. In Sect. 4 and 5 we present numerical examples and discuss some practical choices during the implementation. In Sect. 6 we sum up and discuss the performance of the method.

2. Overview of methods.

Two important questions that characterise any approach to torus construction are: a) how to define and parametrize a model torus, i.e., the map (1), and b) how to fit it in the phase space. In order to guarantee periodicity in the angle coordinates θ, a typical, and natural, torus model involves

Fourier series approximation. The approximated quantities can be simply the phase-space coordinates in the form $q(\theta)$ and $p(\theta)$ (*direct model*), or alternatively, the map (1) can be carried through using the action-angle variables $(\vartheta, \mathcal{J})$ of some known integrable Hamiltonian $\mathcal{H}$ as an intermediate step. In this case, a Fourier series is used to define a canonical transformation $(\theta, J) \leftrightarrow (\vartheta, \mathcal{J})$ by a generating function (*GF model*).

Among the methods of fitting the model torus to the Hamiltonian flow of H, a characterising feature is whether phase-space quantities are sampled along numerically integrated orbits (*trajectory method*), or are they iteratively adjusted on a grid of angles (*iterative method*). Worth noting are also the variables (J, (q, p), or the frequencies ω) which *label*, i.e., define along with the Fourier coefficients, the torus to be modelled.

A diverse selection of methods have been developed, each with their own strengths and weaknesses. Valluri & Merritt [21] give a more thorough review of the subject.

For non-resonant regular orbits, direct trajectory methods, where the fundamental frequencies are determined from the Fourier spectra (e.g., [2, 13, 16]), are a robust choice. Having been refined by many authors, they show good accuracy which, however, decays near resonances, where it takes longer for a sample orbit to fill the torus. In addition, since chaotic orbits do not lie on tori at all, they are a questionable target for all trajectory methods. The hybrid GF-method of Warnock [22] is efficient, but also affected by the limitations above.

For solving the PIP, it is important that tori can be embedded everywhere in the phase space. In this respect, the iterative method of McGill & Binney [14] with its further refinements [1, 10, 12] stands out, because it is unaffected by resonances, and can construct approximations to invariant tori also in phase-space regions where KAM tori do not exist [7, 8]. As a GF-model it has the advantage that the constructed tori are inherently canonical. On the other hand, it relies on the existence of an integrable $\mathcal{H}$ whose tori resemble those of H. The wide class of Stäckel Hamiltonians is available for this purpose, but at the cost of increased analytical and computational effort [12].

The use of simpler, but more restricted $\mathcal{H}$ emphasises a general limitation of GF models; their inability to produce orbits which are arbitrarily thin; i.e., orbits for which $J_\rho \to 0$, where J_ρ is the generalised transversal action, describing the thickness of the orbit in the configuration space [10, 11]. This happens whenever orbits of H and $\mathcal{H}$ with $J_\rho = 0$ do not align in configuration space, which is usually the case. In two-dimensional systems in configuration space, this can sometimes be circumvented by using (rather cumbersome) point transformations [10]. In three dimensions, this becomes much more difficult, and in any case there are tori in systems of all dimensions that cannot be modelled via such transformations. This is an important point when considering the general applicability of a method; indeed, we argue that only the type of method presented in this paper can be called generally applicable.

The main problem with point transformations is that they are only based on the shape of the orbit in configuration space. Thus they separate, e.g., librating and circulation motion explicitly into two topologically different categories though in phase space they share the same topology. Furthermore, the point transformations are used to describe a coordinate system around the orbit $J_\rho = 0$, in which one hopes to be able to describe the orbits $J_\rho \approx 0$ in a consistent way. In general, there is no guarantee that such a system exists; for example, some orbits in a rotating potential

[8] are obviously not describable in this manner. In some cases, the image of the orbit in such a system becomes extremely complicated. Also, in two dimensions, the orbits at $J_\rho = 0$ are closed and thus easily integrated, whereas in three dimensions orbits for which $J_\rho = 0$ describe a surface in configuration space. Constructing this surface to define the desired coordinate system is more difficult than determining a curve.

In this paper, acknowledging the limitations related to trajectory and GF-methods above, we propose a new approach which is direct and iterative. By doing this, we abandon the automatic canonicity of the model, which must be compensated by introducing additional optimisation constraints. A somewhat similar approach was taken by Ratcliff et al. [19], but the present algorithm is less restricted and more robust, since the model is fitted in the least-squares sense, and physically relevant optimisation functions are used. What is more, we introduce a principle for creating tori by letting model tori to adjust themselves by "floating"; i.e., their labels themselves are let to wander locally to some extent, which gives additional flexibility. We illustrate the new method by applying it to selected one- and two-dimensional systems commonly used in galaxy modelling.

3. The torus model.

We represent the model torus as a Fourier series of the Cartesian phase-space coordinates $q \in \mathbb{R}^n$ and velocities $p \in \mathbb{R}^n$;

$$p(\theta) = \sum_{k \in A} \alpha_k e^{i(k\cdot\theta)}, \qquad q(\theta) = \sum_{l \in B} \beta_l e^{i(l\cdot\theta)}, \tag{3}$$

where $A, B \subset \mathbb{Z}^n$ are sets of multi-indices, $\alpha_k, \beta_l \in \mathbb{C}^n$ are Fourier coefficients, and $\theta \in [0, 2\pi)^n$ are canonical angles, conjugate to some actions $J_i \geq 0$, $i = 1, \ldots, n$ which label the model torus. The real-valuedness of p and q dictates that $\alpha_{-k} = \overline{\alpha}_k$ and $\beta_{-l} = \overline{\beta}_l$. The strict consistency of the Cartesian momenta, $p \equiv \dot{q}$, would fix half of the coefficients: $A = B$, $\alpha_k = i(k \cdot \omega)\beta_k$. However, at this point, we wish to retain the maximal flexibility of the model (as a general θ-parameterised surface in the pq-space), and do not enforce this identically.

In order to find the optimal α_k and β_l for a given Hamiltonian $H : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}$, $(q, p) \mapsto H(q, p)$, we build an objective function R which is a weighted sum

$$R(\theta) = \sum_i \lambda_i \|\mathcal{E}_i(\theta)\|^2 \tag{4}$$

of error functions $\mathcal{E}_i$ that vanish on a KAM torus of H. The norm is Euclidean, and $\lambda_i > 0$. We will minimise R on a grid of angles $\theta_{(m)}$, $m = 1, \ldots, M$ with respect to the Fourier coefficients α_k, β_l. Obviously, success in torus construction requires that the error functions $\mathcal{E}_i$ are chosen wisely. This is our goal in the following.

Hamilton's equations of motion under H are

$$\dot{p} = -\partial_q H, \quad \dot{q} = \partial_p H, \tag{5}$$

and supposing that (θ, J) are action-angle variables of H, we also have

$$\dot{J} = 0, \quad \dot{\theta} = \partial_J H = \omega. \tag{6}$$

Applying the chain rule to the time derivatives of the model (3), inserting (6) and equating to (5) yields

$$\frac{\partial p}{\partial \theta}\omega = -\frac{\partial H}{\partial q}, \quad \frac{\partial q}{\partial \theta}\omega = \frac{\partial H}{\partial p} \tag{7}$$

which identifies the vector field produced by the model with the Hamiltonian vector field of H. Assuming that H is differentiable, the partial derivatives in Eq. (7) can be evaluated for all θ. We define two error functions,

$$\mathcal{E}_1(\theta) \coloneqq \frac{\partial p}{\partial \theta}\omega + \frac{\partial H}{\partial q}, \quad \mathcal{E}_2(\theta) \coloneqq \frac{\partial q}{\partial \theta}\omega - \frac{\partial H}{\partial p}, \tag{8}$$

which penalise for any deviation from the Hamiltonian flow of H. Obviously, in order to use Eq. (8) we need the values of the frequencies ω. A simple solution would be to use ω for labelling the model torus, i.e., to use ω as constants. However, sets of valid frequencies can be difficult to sample, and they may overlap for different orbit families, which makes J the preferred label.

Actually, we can avoid labelling the model torus (beforehand) altogether, if we dictate that ω is a least-squares estimator, obtained by fitting the $2n$ equations (7) at all of the grid points. More specifically, we have an overdetermined linear system of the form $A\omega = b$, where A is an $(2nM \times n)$-matrix obtained by concatenating the $(n \times n)$-matrices $\partial p/\partial \theta_{(m)}$ and $\partial q/\partial \theta_{(m)}$ for all $m = 1, \ldots, M$. The solution is given by the normal equations;

$$\omega = (A^T A)^{-1} A^T b. \tag{9}$$

This complicates the dependence of $\mathcal{E}_1$ and $\mathcal{E}_2$ on the Fourier coefficients somewhat, but proved to be a numerically robust solution in our experiments.

Convergence in minimisation towards a KAM torus of H is possible with just the two error functions (8). This principle was employed by Ratcliff et al. [19], with the aim of having the two vanish at a given number of phase-space points. This, however, causes uncontrollable fluctuations on the rest of the constructed torus when it does not approximate a KAM one of H. Our reconstruction is based on the geometric principle of least-squares fitting a number of error functions $\mathcal{E}_i$ that are equivalent on KAM tori (subsets of them are necessary and sufficient), but different elsewhere. In other words, we construct the desired manifolds by regularisation and global smoothing.

From Hamilton's equations of motion (6) we can easily derive

$$\mathcal{E}_3(\theta) \coloneqq \frac{\mathrm{d}J}{\mathrm{d}t} = \frac{\partial H(q(\theta), p(\theta))}{\partial \theta} = \frac{\partial H}{\partial q}\frac{\partial q}{\partial \theta} + \frac{\partial H}{\partial p}\frac{\partial p}{\partial \theta} \tag{10}$$

which gives additional support to, and falls into the same category as, $\mathcal{E}_1$ and $\mathcal{E}_2$. In numerical exercises, where Eq. (9) was used, we found that the inclusion of $\mathcal{E}_3$ improved our results noticeably.

Further regularisation is probably necessary, at least when constructing tori in non-regular regions of phase space of H. A useful error function is obtained by dictating that the value of the Hamiltonian H should be constant on the minimisation grid;

$$\mathcal{E}_4(\theta) \coloneqq H(\theta) - \bar{H}, \tag{11}$$

where $\bar{H}$ is the arithmetic mean of H over the grid. This also yields the local solution of PIP: we define $H_0 \coloneqq \bar{H}$ on the constructed torus. The global interpolation scheme over the tori, must, of course, fulfill $\omega(J) = \partial H_0/\partial J$ everywhere.

Finally, we introduce an error function which ensures that the model converges towards a specific torus with given values of J. For this purpose, we need to derive

expressions for the model actions. By definition,

$$J_h = \frac{1}{2\pi} \oint_{\gamma_h} p \cdot \mathrm{d}q, \quad h = 1, \ldots, n,$$

where γ_h is a closed path that cannot be continuously deformed into a point. For a model of the form $q(\theta)$ we have

$$\mathrm{d}q = \sum_{j=1}^{n} \frac{\partial q}{\partial \theta_j} \mathrm{d}\theta_j,$$

and if we select the path γ_h in such a way that $\mathrm{d}\theta_j = 0$ for $j \neq h$, we can write

$$J_h = \frac{1}{2\pi} \sum_{j=1}^{n} \int_0^{2\pi} p_j \frac{\partial q_j}{\partial \theta_h} \mathrm{d}\theta_h.$$

By inserting the model equations (3) and integrating, we obtain

$$J_h = \sum_{j=1}^{n} \sum_{(k,l) \in D_h} i k_h \alpha_{j,k} \beta_{j,l} e^{i[(k+l)\cdot\theta]}, \tag{12}$$

where $D_h = \{(k,l) \in A \times B : l_h = -k_h\}$, and $\alpha_{j,k}$ and $\beta_{j,l}$ are the j:th components of the corresponding Fourier coefficients. Note that J_h does not depend on θ_h. On a KAM torus, the dependence on θ_j, $j \neq h$ should vanish as well. In order to enforce this, and to punish for any deviation from a given set of actions $J = \bar{J}$, we introduce an error function

$$\mathcal{E}_5(\theta) \coloneqq J(\theta) - \bar{J}. \tag{13}$$

As our notation suggests, $\bar{J}$ could also be computed as an arithmetic mean over the grid, supplementing the label-free version of the torus construction algorithm.

4. **Numerical example: isochrone potential.** Let us first check how the method reconstructs simple integrable Hamiltonians. The harmonic oscillator is trivial, since it is given directly in the Fourier form of Eq. (3) with only first-order terms. As a more interesting sanity check for the presented torus model and minimisation criteria, we apply them to the isochrone Hamiltonian [6] with a single degree of freedom;

$$H(q,p) = \frac{1}{2}p^2 - \frac{c_1}{c_2 + \sqrt{c_2^2 + q^2}},$$

where $q, p \in \mathbb{R}$, and the parameters have values $c_1 = 1$ and $c_2 = 0.15$. By fixing the origin of θ to $q = 0$, and acknowledging the symmetries of the system, we express the torus model (3) in trigonometric form;

$$p(\theta) = \sum_{k \in X} a_k \cos(k\theta), \qquad q(\theta) = \sum_{l \in Y} d_l \sin(l\theta),$$

where $X = Y = \{1, 3, 5, 7, \ldots, N-1\}$. Since the system is one-dimensional and integrable, we can, in this case, safely label the model tori by their frequencies ω, which simplifies the error functions (8). We have

$$\begin{aligned}
\mathcal{E}_1(\theta) &= \omega p'(\theta) + \frac{c_1 q(\theta)}{\sqrt{c_2^2 + q^2(\theta)}\left(c_2 + \sqrt{c_2^2 + q^2(\theta)}\right)}, \\
\mathcal{E}_2(\theta) &= \omega q'(\theta) - p(\theta),
\end{aligned}$$

where

$$p'(\theta) = -\sum_{k\in X} k a_k \sin(k\theta), \qquad q'(\theta) = \sum_{l\in Y} l d_l \cos(l\theta).$$

The rest of the error functions $\mathcal{E}_i$, $i = 3, 4, 5$ will not be used in this example. We minimise the objective function (4) (with $\lambda_1 = \lambda_2 = 1$) using the Levenberg-Marquardt algorithm (LMA). This requires the calculation of the partial derivatives $\partial\mathcal{E}_i/\partial a_k$ and $\partial\mathcal{E}_i/\partial d_l$, $i = 1, 2$, which is straightforward using the equations above.

With the purpose of studying the convergence radius and accuracy of the method, we ran the LMA several times, while varying the number of coefficients $N = 16, 32, 48, \ldots, 512$ and the frequency $\omega = 0.1, 0.15, 0.2, \ldots, 2.0$. The angle grid contained $M = 1024$ equally spaced values (matching the Nyquist rate for $N = 512$), but due to the symmetry of the model, only half of them were actually needed and used. As an initial guess for the model torus we set all the coefficients to zero except $a_1 = d_1 = 1$, which corresponds to a unit circle in the qp-plane (i.e., a harmonic oscillator). For each pair (N, ω), after the LMA was well converged, we evaluated the standard deviation σ of the Hamiltonian $H(q(\theta), p(\theta))$ over the grid points. Figure 1 shows a contour plot of the results.

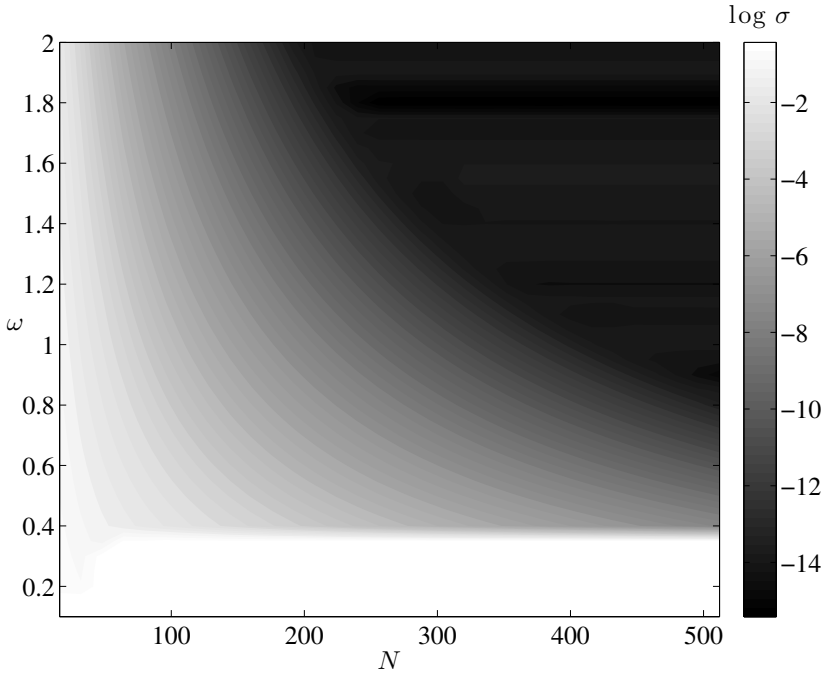

FIGURE 1. Standard deviation σ of the one-dimensional isochrone Hamiltonian function values on the θ-grid after running the torus embedding LMA.

For all $\omega \geq 0.4$ the LMA converged, and H attained a constant value throughout the grid, within the accuracy seemingly bestowed by N and also ω. In the upper right portion of the figure, the values of σ are saturated due to the finite computational precision (64-bit floating point). For the isochrone, we know the relation

$H = -(2c_1\omega)^{2/3}/2$, and hence, we could verify that Fig. 1 also closely represents the deviation from the analytical values $H(\omega)$. For $\omega < 0.4$ there was no proper convergence. However, this could be remedied by compensating the increased amplitude of the motion in the initial guess ($\omega \to 0 \implies J \to \infty$). For example, with an initial guess $q(\theta) = 2\sin\theta$, $p(\theta) = 2\cos\theta$ the LMA converged nicely for $\omega = 0.2, \ldots, 0.4$. We expect that a decent initial guess is increasingly important for more complex systems.

5. **Numerical example: logarithmic potential and the PPS.** We raise the degrees of freedom, and formulate the torus modelling algorithm for planar systems of the form

$$H(q,p) = \frac{1}{2}p \cdot p + \Phi(q),$$

where $q, p \in \mathbb{R}^2$ and $\Phi : \mathbb{R}^2 \to \mathbb{R}$ is a potential function. Depending on Φ, there may be structural features in phase space, such as differentiation into orbit families, and stochastic regions surrounding resonant islands (for near-integrable potentials). In order to obtain a basic understanding how the torus algorithm copes with these features, we apply it to two selected potentials. First, we have the near-integrable logarithmic potential [20, 4]

$$\Phi(q) = \frac{1}{2}\ln\left(q_1^2 + \frac{q_2^2}{c_1^2} + c_2^2\right) \tag{14}$$

where $c_1 = 0.9$ and $c_2 = 1$. The phase space of the logarithmic system is dominated by KAM tori, and with the selected parameter values, most of the orbits are either boxes (i.e. butterflies) or loops [15].

For comparison, as an integrable example, we use a Stäckel potential in elliptic coordinates u;

$$\Phi(u) = -\frac{f(u_1) - f(u_2)}{u_1 - u_2}, \tag{15}$$

where f is a smooth function. The coordinate transformation $q \mapsto u$ is given by

$$u_1 + u_2 = -c_1 - c_2 + q_1^2 + q_2^2, \qquad u_1 u_2 = c_1 c_2 - c_2 q_1^2 - c_1 q_2^2,$$

where $c_1 < c_2$ are parameters. By selecting $u_2 \le u_1$, we have $-c_2 \le u_2 \le -c_1 \le u_1$. With the choice

$$f(u_1) = -2\pi(u_1 + c_1)c_2c_3\sqrt{\frac{-c_1}{u_1 + c_1}}\arctan\sqrt{\frac{u_1 + c_1}{-c_1}},$$
$$f(u_2) = -2\pi(u_2 + c_1)c_2c_3\sqrt{\frac{c_1}{u_2 + c_1}}\operatorname{artanh}\sqrt{\frac{u_2 + c_1}{c_1}},$$

the potential (15) becomes the perfect prolate spheroid (PPS), a special case of the perfect ellipsoid [5]. We set the parameter values to $c_1 = -1$, $c_2 = -0.25$, and $c_3 = 1$. In the q_1q_2-plane, the PPS hosts the same major orbit families, the boxes and loops, as the logarithmic potential (14). In fact, the phase-space structures of the two systems are very similar, and in a side-by-side comparison of the corresponding torus models, integrability of the PPS supports the analysis on the logarithmic side as well.

Since we expect to model only KAM tori, our main concern is how the orbit family, or the shape of an orbit within the family, affects the outcome of the torus modelling. Compared to the one-dimension case, a big difference is that we now use the actions J to label the tori. There are two reasons for not using ω. First, for the

PPS (and thus possibly for the logarithmic system as well), the values of ω overlap for boxes and loops (Fig. 2 and 3). Second, in some of our experiments while using ω as a label, the LMA converged to a trivial 1-torus ($J_1 = 0$ or $J_2 = 0$).

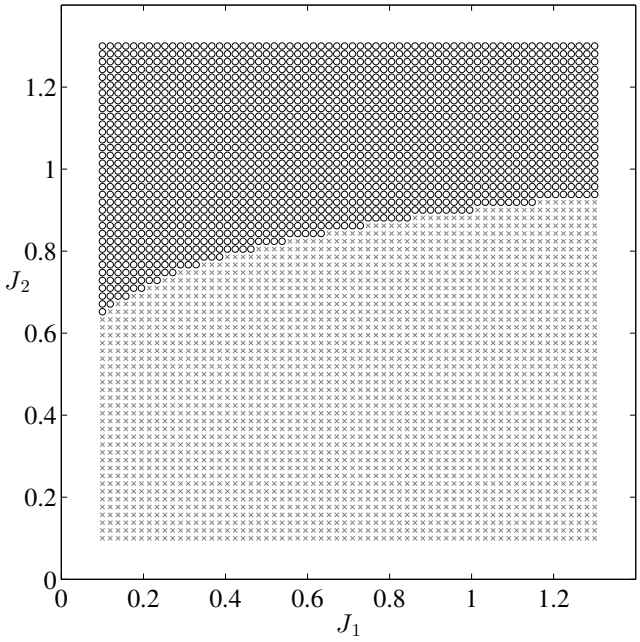

Figure 2. The loops (black circles) and boxes (gray crosses) in the action space of the PPS.

A general trigonometric version of the torus model (3) is now

$$\begin{aligned} p(\theta) &= \sum_{k\in X} a_k \cos(k\cdot\theta) + b_k \sin(k\cdot\theta), \\ q(\theta) &= \sum_{l\in Y} c_l \cos(l\cdot\theta) + d_l \sin(l\cdot\theta), \end{aligned} \tag{16}$$

where $q, p \in \mathbb{R}^2$, $a_k, b_k, c_l, d_l \in \mathbb{R}^2$, $\theta \in [0, 2\pi)^2$, and $X, Y \subset \mathbb{Z}^2$ are sets which do not contain opposite indices. Due to symmetries and other characteristics of a given orbit (family), some terms in Eq. (16) are always zero. Significant savings in computation time of the LMA are obtained, if the corresponding indices are identified and omitted from X and Y. For the integrable PPS there is a shortcut for doing this numerically. Namely, since we know the transformation $(\theta, J) \mapsto (q, p)$, we can compute sample points $q(\theta)$ and $p(\theta)$ on an even θ-grid, and perform a discrete Fourier transformation (DFT) which yields a representation exactly of the form of our model. The computation is relatively fast (with optimised FFT-routines), and the DFT orbit can also be used to verify the result of the LMA. For the PPS box orbits, we found the nonzero coefficients to be

$$\begin{aligned} &a_{1,k},\ \text{if } k = (\text{odd, even}), \qquad &&d_{1,l},\ \text{if } l = (\text{odd, even}), \\ &a_{2,k},\ \text{if } k = (\text{even, odd}), \qquad &&d_{2,l},\ \text{if } l = (\text{even, odd}), \end{aligned} \tag{17}$$

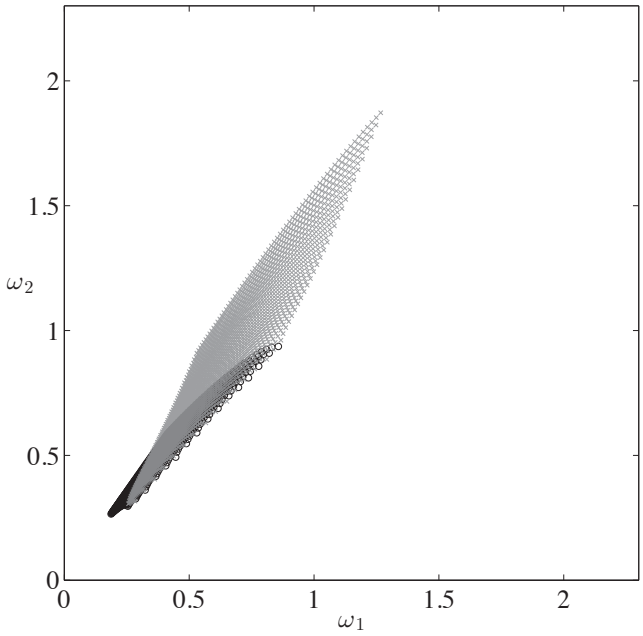

FIGURE 3. The orbits of Fig. 2 mapped into the PPS frequency space.

and for the loops;

$$\begin{aligned} &b_{1,k},\ \text{if } k = (\text{even, odd}), \qquad &&c_{1,l},\ \text{if } l = (\text{even, odd}),\\ &a_{2,k},\ \text{if } k = (\text{even, odd}), \qquad &&d_{2,l},\ \text{if } l = (\text{even, odd}). \end{aligned} \tag{18}$$

We define two separate initial guesses for the model (16); one for box orbits, and the other for loops, containing only coefficients of the form (17) and (18), respectively. Because of the common orbit families, these initial guesses work for both the PPS and the logarithmic potential.

Let $N \in \mathbb{N}$ limit the range of the indices such that $|k_i| \leq N$, $|l_i| \leq N$, $i = 1, 2$. By setting $N = 16$, we end up with a total of 544 coefficients per orbit which are initially set to zero except the ones appearing below. As an initial box orbit, we use independent harmonic oscillators in both coordinate directions;

$$q(\theta) = \begin{pmatrix} \sin\theta_1 \\ \sin\theta_2 \end{pmatrix}, \qquad p(\theta) = \begin{pmatrix} \cos\theta_1 \\ \cos\theta_2 \end{pmatrix}, \tag{19}$$

which corresponds to $\omega = (1, 1)$. For a loop orbit we must include more nonzero terms, in order to have the initial $J_1 > 0$ in Eq. (12). After some experiments, we found that

$$q(\theta) = \begin{pmatrix} \cos\theta_2 + \dfrac{1}{20}\cos(2\theta_1 + \theta_2) - \dfrac{1}{2}\cos(-2\theta_1 + \theta_2) \\ \dfrac{3}{2}\sin\theta_2 + \dfrac{1}{10}\sin(2\theta_1 + \theta_2) - \dfrac{1}{2}\sin(-2\theta_1 + \theta_2) \end{pmatrix}, \tag{20}$$

and the $p(\theta)$ that follows from selecting $\omega = (\frac{1}{2}, \frac{1}{2})$ are adequate.

Using the least-squares frequencies (9), the error functions $\mathcal{E}_1$ and $\mathcal{E}_2$ in Eq. (8) for all grid points are given by $f(\theta_{(m)}, m = 1, \ldots, M) = A\omega - b$. Differentiation with respect to the model coefficients gives

$$\frac{\partial f}{\partial x} = A(A^T A)^{-1} \frac{\partial A^T}{\partial x}(b - A\omega) + (A(A^T A)^{-1} A^T - I)\left(\frac{\partial b}{\partial x} - \frac{\partial A}{\partial x}\omega\right), \tag{21}$$

where $x = a_k, b_k, c_l, d_l$. Since the partial derivatives required for $\partial\mathcal{E}_3/\partial x$ and $\partial\mathcal{E}_4/\partial x$ already appear in Eq. (21), there is no trouble in adding $\mathcal{E}_3$ and $\mathcal{E}_4$ to the objective function (4). In order to use $\mathcal{E}_5$ we must evaluate the actions (12) for the trigonometric version of the model. Taking care of the indexing, we have, for any n,

$$\begin{aligned} J_h = -\frac{1}{2}\sum_{j=1}^{n}\sum_{(k,l)\in Z_h} k_h \Big\{ &(a_{j,k}d_{j,l} + b_{j,k}c_{j,l})\cos\left[(k+l)\cdot\theta\right] \\ &+(b_{j,k}d_{j,l} - a_{j,k}c_{j,l})\sin\left[(k+l)\cdot\theta\right]\Big\}, \end{aligned} \tag{22}$$

where $Z_h = \{(k,l) \in X \times Y \cup X \times -Y : l_h = -k_h\}$.

We set the θ-grid to be a 32×32 square lattice. This fulfils the requirements of the Petersen-Middleton theorem [17] for the chosen number of model coefficients (i.e., there should be no aliasing). However, due to model symmetries with respect to both coordinate axes, we can omit everything but the values $\theta \in [0, \pi)^2$. This leaves us with only 256 grid points per orbit. Nevertheless, compared to the one-dimensional case, the computational effort to construct a torus is now far greater. Therefore, we do not try to optimise near to machine precision in the LMA. Instead, we set modest goals for the objective function value and concentrate on finding out whether the algorithm converges towards the correct solution or not.

5.1. **Unlabelled torus construction.** In our first set of experiments, we use an objective function (4) which is unlabelled, consisting of the error functions $\mathcal{E}_i$, $i = 1, \ldots, 4$ with the weights λ_i set to unity. We run the LMA twice for each system (the PPS and logarithmic); first using the fixed initial guess for box orbits (19) and then using the one for loops (20). In all four cases, the LMA converges fast during the first few iteration steps, before reaching a plateau of slow convergence. During the phase of slow convergence, there may still be apparent changes in the model parameters. We interpret this behaviour being caused by the algorithm first finding a valley of KAM tori, and then advancing along its bottom. Since we have not specified a label, any KAM torus will suffice, and we tune the LMA to stop as soon as the bottom of the valley is reached. For comparing LMA results we again use σ; the standard deviation of the corresponding Hamiltonian function H over the θ-grid. Table 1 lists the final values of the actions, frequencies, and the values of σ.

It seems that the logarithmic tori are easier to construct than the PPS ones. For both systems, the constructed boxes are more accurate than the corresponding loops. It can be verified that the actions and frequencies in Table 1 all correspond to invariant tori. For the PPS we can do this analytically (cf. Fig. 2), and for both systems we can visually compare constructed orbits to numerically integrated ones, when the orbits share the same initial values. Fig. 4 shows the Poincaré section $y = 0$, $\dot{y} > 0$ for the orbits in Table 1. The common initial point is $\theta = (0, \pi/2)$ on the constructed torus. The numerical integration of Hamilton's equations of motion

TABLE 1. Initially unlabelled constructed tori for the PPS and the logarithmic potential.

		J	ω	$\log\sigma$
PPS	box	$(0.19, 0.34)$	$(0.97, 1.30)$	1×10^{-5}
	loop	$(0.14, 1.23)$	$(0.43, 0.60)$	8×10^{-5}
logarithmic	box	$(0.16, 0.22)$	$(0.78, 0.87)$	6×10^{-7}
	loop	$(0.11, 0.76)$	$(0.58, 0.67)$	2×10^{-6}

in Cartesian coordinates is done using the Gragg-Bulirsch-Stoer (GBS) method with high accuracy settings (the standard deviation of H along the integrated orbits is below 10^{-12}). The plotted orbits coincide in all four cases. Hence, the tori are

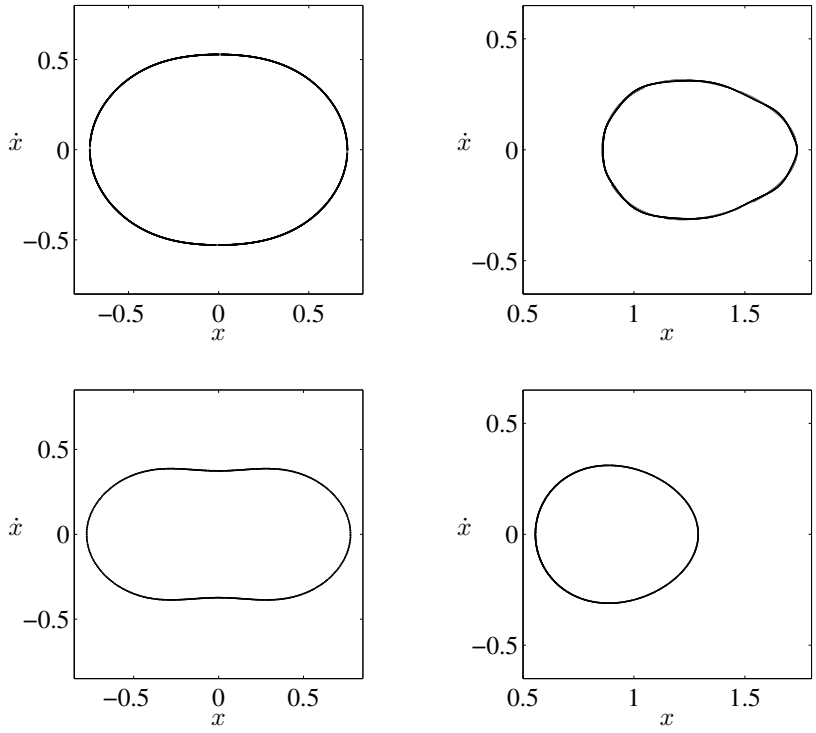

FIGURE 4. Poincaré sections $y = 0$, $\dot{y} > 0$ for the unlabelled torus construction, superimposed over numerically integrated reference orbits (gray): the PPS box (upper left) and loop (upper right), and logarithmic box (lower left) and loop (lower right).

indeed close to invariant.

For the PPS, increasing the number terms by setting $N = 24$ (total of 1200 coefficients), and resizing the θ-grid accordingly, resulted in less than an order of magnitude gain in the final σ. For our purposes, this was insignificant compared to the increased computation time, and we decided to stick with $N = 16$ in the following.

5.2. **Labels from an action grid.** Our next goal is to form an overall picture of how the torus construction performs when the actions J are given as a label. An obvious problem in this approach is that we should know the orbit family in advance in order to supply the proper initial torus to the algorithm. Also, based on the isochrone experiment, the LMA may fail, if the size of the initial torus is far from correct. However, in the case of the integrable PPS, we have an opportunity to cheat, and to bypass these difficulties by using an initial torus given by the DFT. For a given action label, the set of non-zero Fourier coefficients from the DFT automatically determines the orbit family, and the truncated series is an excellent initial guess for the torus algorithm. What we obtain is a best-case scenario for the chosen size of the model (N, M). Figure 5 shows the results for action labels matching the range in Fig. 2. Clearly, the accuracy of the fit decreases as the tori

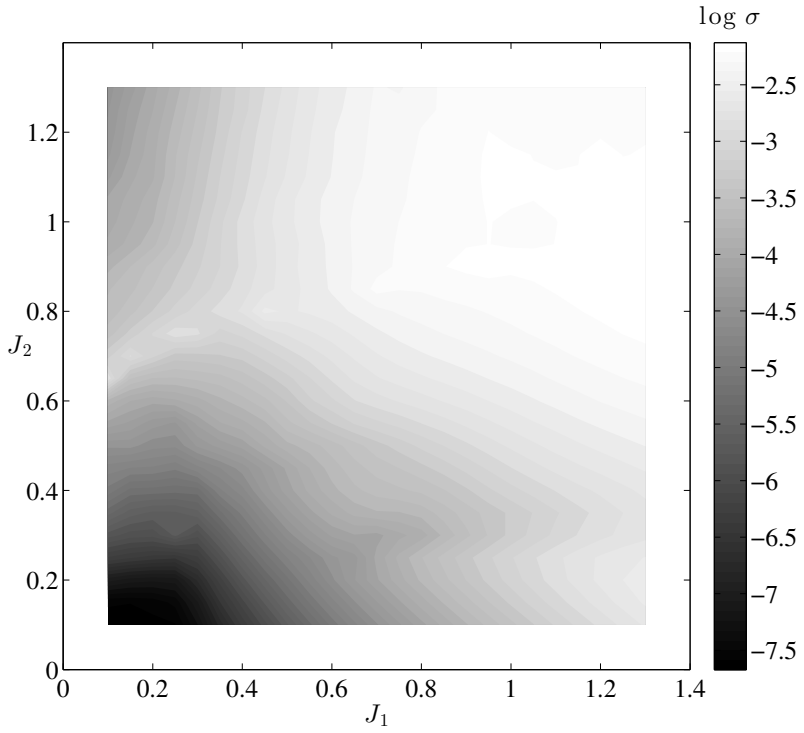

FIGURE 5. The final standard deviation σ of the PPS Hamiltonian H on the θ-grid using initial guesses from the DFT.

get thicker and also near the separatrix of the box and loop orbits. The tori found by the unlabelled torus construction (Table 1) are located in areas of low σ, which seems natural.

Next, we shall take a brute force approach which does not rely on integrability, and run the LMA twice for each action label, first using the initial box orbit (19), and then the initial loop (20). From now on, the labels are taken from a rectangular grid of actions, consisting of the values $J_1 = 0, 0.05, 0.1, \ldots, 1.3$ and $J_2 = 0, 0.05, 0.1, \ldots, 1.3$ (note the inclusion of orbits with zero thickness). The action grid contains both boxes and loops for the PPS, and we shall see that this is

also the case for the logarithmic potential. The same grid is used for both systems. The objective function (4) is now fully fledged, including all $\mathcal{E}_i$, $i = 1, \ldots, 5$, and the corresponding weights λ_i set to unity. The results for each potential-orbit-family pair are gathered in Fig. 6.

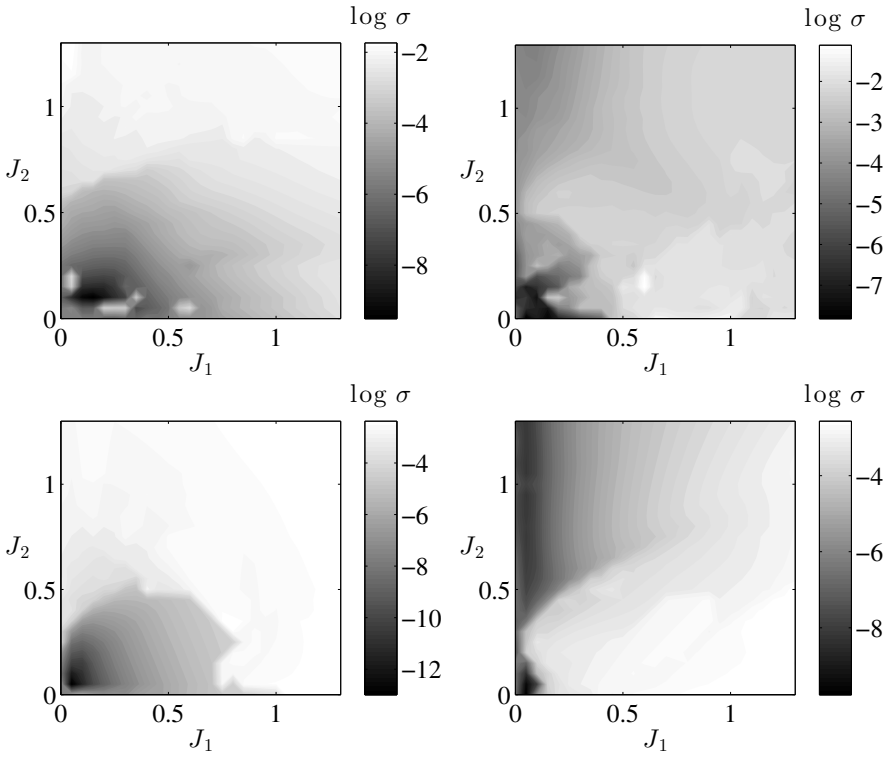

FIGURE 6. The final standard deviation σ of the Hamiltonian H on the θ-grid for the labelled torus construction using the fixed initial guesses: the PPS box (upper left) and loop (upper right), and logarithmic box (lower left) and loop (lower right).

What we would like to see, is that there would be a clear threshold in σ, differentiating the regions where the initial orbit family is correct and where it is not. This would give us a tool for determining the orbit family by trial and error. Unfortunately, at least alone, σ is inadequate for such a purpose. However, most of the structures in Fig. 5 can be identified in Fig. 6, and one can clearly visualise the separatrix of boxes and loops for the logarithmic potential.

It seems that in regions where the initial orbit family is correct, the algorithm is not overly sensitive to the choice of the initial torus. There are a few isolated high values of σ for the PPS boxes, and a systematic change for the logarithmic boxes as J_1 increases (cf. Fig. 1). These (local minima) can be avoided by improving the initial guess, as we shall see later.

Although not emphasised in Fig. 6, the torus construction can run into various problems near the separatrices. First of all, there are some tori which attain low values of σ despite being stuck in a local minimum. Fortunately, these can be

identified by having an inconsistent set of Fourier coefficients in such a way that the values $|\alpha_k - i(k\cdot\omega)\beta_k|$, $k \in A \cap B$ in (3) are large. Then, even if the LMA is heading towards the correct solution, it may still suffer from slow convergence. Finally, if a good fit is obtained, the orbit family may still be incorrect, since σ is transparent in this respect (Fig. 7). On the other hand, the algorithm does not

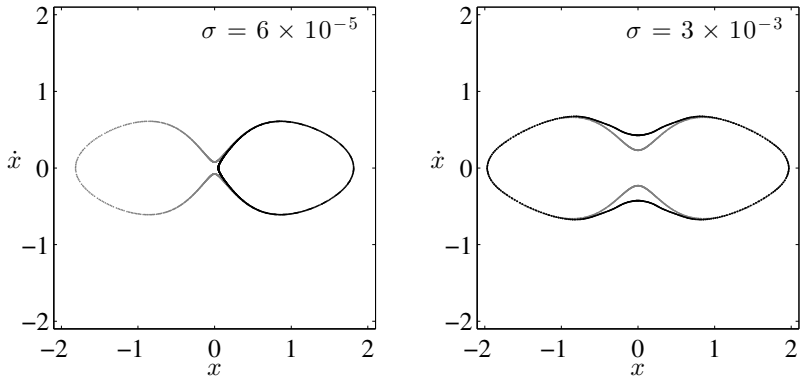

FIGURE 7. Poincaré sections $y = 0$, $\dot{y} > 0$ for selected constructed tori near the separatrix, superimposed over numerically integrated reference orbits (gray): the logarithmic loop $J_1 = 0.5, J_2 = 0.65$ (left), and the PPS box $J_1 = 0.65, J_2 = 0.75$ (right).

seem to be bothered by thin orbits; even the ones with zero thickness are modelled without a notable drop in accuracy (Fig. 8; the box orbit on the right is shown for symmetry, though it is one of the rare cases when GF-methods work for thin orbits without point transformations).

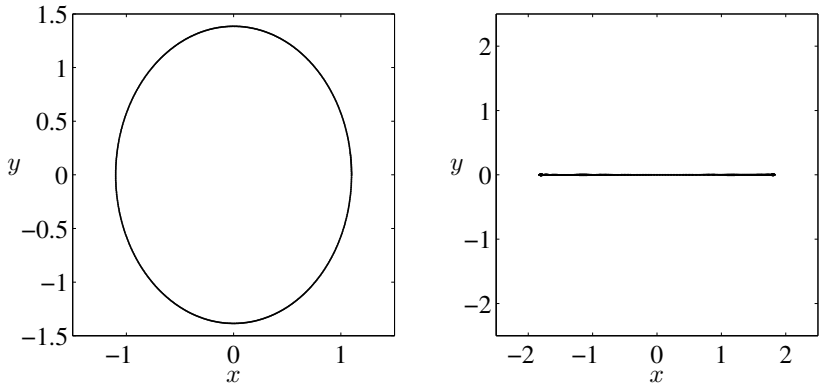

FIGURE 8. Constructed thin orbits in the xy-plane: the logarithmic loop $J_1 = 0, J_2 = 1$ (left), and box $J_1 = 1, J_2 = 0$ (right), superimposed over numerically integrated orbits (gray).

Curiously, it seems that the model for loop orbits (20) is able to mimic a box orbit for small values of the actions, by setting a large amount of its coefficients to zero. Although σ for these orbits is relatively low, they are not close to the corresponding solutions (equal action labels) obtained by using the box model (19). We are not too concerned about this anomaly, though, since it does not play a role, if we, instead of blindly going through the whole action grid, use the following algorithm to construct tori only for selected grid points.

5.3. **Action-grid probing.** Based on the the experiments above, we propose a scheme for probing the action grid, and isolating the regions where the torus construction best succeeds. The idea is to start from a point in the action grid which is closest to the actions obtained from the unlabelled torus construction. If the construction for the starting point is successful, we use it as the initial torus for all the adjacent grid points. When this procedure is iteratively repeated, we obtain an expanding region of successfully constructed tori in the action grid. Besides avoiding local minima, this technique enhances the rate of convergence in the LMA. A detailed description of the algorithm is given in Appendix A.

The algorithm must repeatedly evaluate whether a constructed torus is good or not. The choice is simply based on the final value of the objective function. If the value exceeds a certain threshold, the torus is rejected. In the following, we use a safe value of 10^{-6}, but this is subject to fine tuning. In order to rule out the specific local minima near the separatrices, we modify the objective function by adding the term

$$R_0 = \rho \sum_{k \in C} |\alpha_k - i(k \cdot \omega)\beta_k|^2, \tag{23}$$

where $C = A \cap B$, and ρ is a weighting parameter. We set $\rho = 0.01/\#C$.

The results of action-grid probing, using the tori in Table 1 as initial values, are displayed in Fig. 9-12. For the PPS, the results in Fig. 2-3 are partly reconstructed, allowing a direct comparison. Note, however, that the thinnest orbits are only present in the constructed case.

The points in the action-space map to frequency space as follows. In Fig. 10 and 12 the cusps pointing northeast correspond to the origin $J_1 = J_2 = 0$. The points $J_2 = 0$ form the northern facets. The eastern facets correspond to the box orbits with $J_1 = 0$. The southern facet in Fig. 3 (and partly in Fig. 12) corresponds to the separatrix of the boxes and loops. For both the PPS and logarithmic potential, a few box orbits with zero thickness near the separatrix seem to have slightly incorrect values of ω. Otherwise, the results seem consistent.

6. **Discussion.** We have introduced the Poincaré inverse problem and proposed a new method towards its solution via torus construction. We have applied the method to the isochrone, PPS, and logarithmic systems. For the one-dimensional isochrone, the torus construction trivialises, but the results still prove the point that the LMA may not converge, if the initial torus is far away from the labelled KAM torus. Also, it is evident that the number of Fourier terms, needed to maintain a certain level of accuracy, increases as the tori get thicker ($\|J\|$ increases). These observations can be made for the PPS and logarithmic systems as well. In addition, in the two-dimensional cases, torus construction becomes more difficult near the separatrix between the orbit families. However, the thin orbits pose no problem, which may be a major selling point for our method.

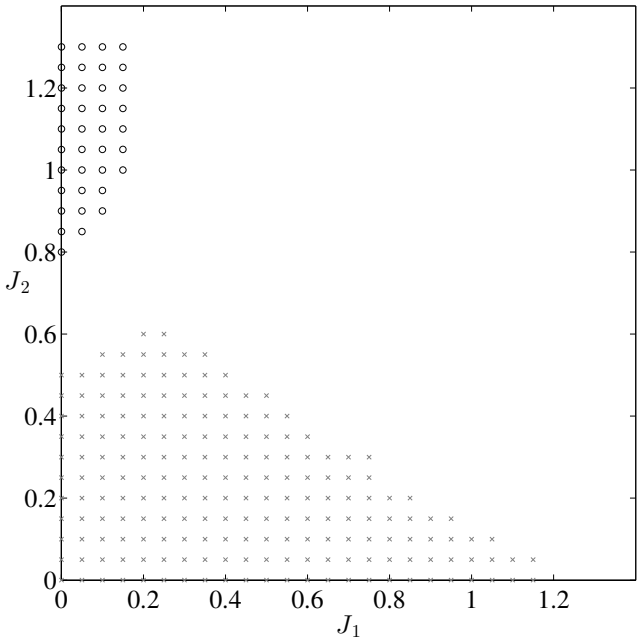

FIGURE 9. The PPS loops (black circles) and boxes (gray crosses) constructed by the action-grid probing algorithm.

Especially for the PPS, the size of the torus model (set according to our computational budget) limits the accuracy of the results. Nevertheless, by using the action-grid probing algorithm, one is able to cherry-pick the most accurately constructed tori. For the specific goal of constructing KAM tori (boxes and loops) for the near-integrable logarithmic potential, the algorithm works quite well. This culminates in Fig. 11-12 where we provide a (partial) distribution of orbit families in the action and frequency spaces of a system which is inherently non-integrable. However, a question yet unanswered is how the modelling succeeds for the boxlet orbits; the bananas, fish, and pretzels (see, e.g., [15]), which appear when $c_2 \to 0$ in the logarithmic potential (14). Our results are similar to the action diagrams of Binney & Spergel [3] which were obtained with slightly different potential parameters.

If increasing the size of the model is not an option, there are a couple of ideas that one could try to improve the quality of the constructed tori. First, the Fourier coefficients could be included according to $\|k\| \leq N$, instead of $|k_i| \leq N$, $i = 1, \ldots, n$. This, in combination with a hexagonal angle grid, would provide a more natural sampling scheme and could have a positive effect on the results. Second, the choice that the elementary phase-space variables q and p are Cartesian were made purely based on convenience; in the torus model (1) they can be defined to be any set of variables, as long as those uniquely define a point in the phase space. The Fourier series will certainly be shorter, if the geometry of the system matches the selected variables. However, one should note that angle-like variables (e.g., among polar, cylindrical, or spherical coordinates) are not suitable to be modelled

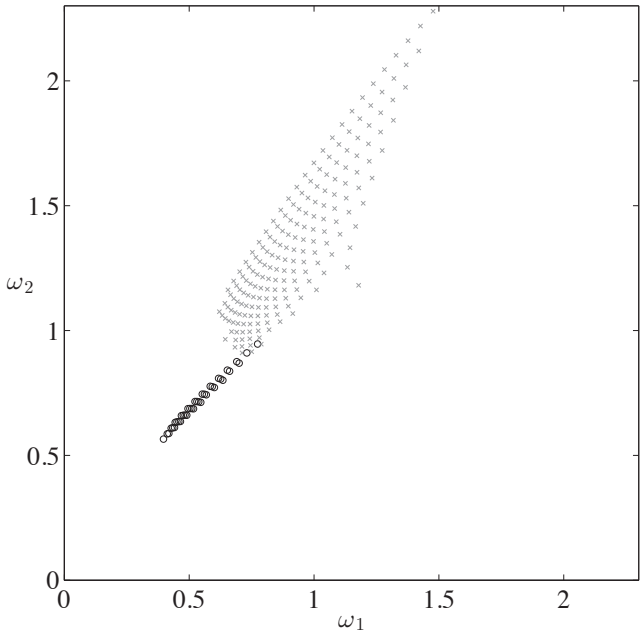

FIGURE 10. The constructed PPS frequencies, corresponding to the orbits in Fig. 9.

using Fourier series because of their discontinuous nature. A trigonometric function of such a variable is obviously fine, but introduces ambiguity. Also, in curvilinear coordinates the computation of the partial derivatives needed for the LMA becomes more cumbersome. In our two-dimensional experiments, the Cartesian model was better suited for the logarithmic system. This might be explained by the fact that the PPS separates in elliptic, not in Cartesian, coordinates, although the orbit shapes are similar in the two systems.

How does the new method with a direct torus model compare against iterative GF-methods? Its biggest advantage is undoubtedly flexibility; the method performs well for thin orbits and can also handle orbits near the separatrix. As a definite downside, it seems that increasing the size of the direct model is more expensive in computation time. This is understandable, since the LMA for a GF-method has to cope with only one Fourier series, and its objective function is simpler. Also, in the direct model, the built-in least squares solution (9) gets computationally heavy for large θ-grids. Finding a suitable initial torus is important in both approaches; with the GF-model, we need to find a nearby integrable toy torus in order to initialise the construction process. This can be nontrivial. For the direct model, in order to avoid local minima, an equally nontrivial set of Fourier coefficients is required. This, however, can be obtained in various ways. The GF-model may have an edge when the full solution $J \mapsto H_0(J)$ to the PIP is sought, since the Poincaré integral conditions are automatically satisfied on the interpolating tori; the direct model may require a denser grid of constructed tori.

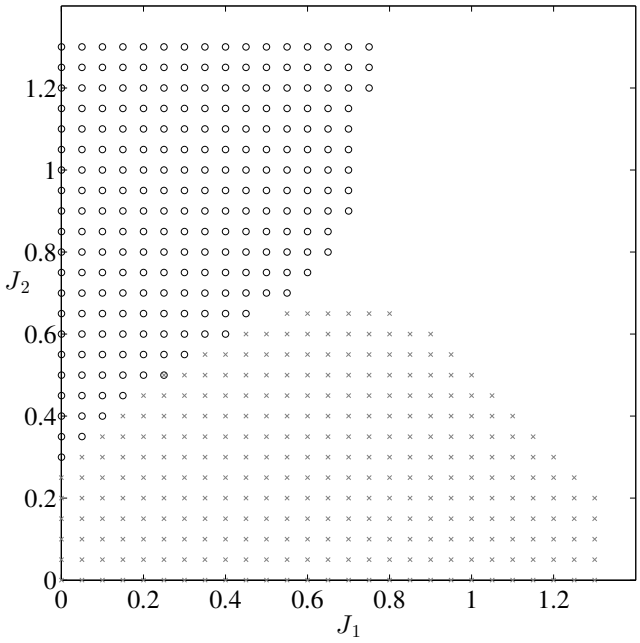

FIGURE 11. The logarithmic loops (black circles) and boxes (gray crosses) constructed by the action-grid probing algorithm.

In the big picture, the algorithm presented in this paper complements the arsenal of torus construction methods. Within our classification, all four combinations of iterative/trajectory and direct/GF methods, have been tried. This, of course, does not mean that everything is invented; tori are geometrical objects in phase-space, and if viewed purely as such, they can be represented in many ways. One possibility could be to compute a best-fit tessellation and an angle-coordinate system for a bunch of numerically integrated orbit strips. As a simple example, consider Poincaré sections, cut along consecutive azimuthal angles, of a two-dimensional loop orbit with a constant value of H, and visualise the torus model. As a benefit, this kind of approach would not be tied to Hamiltonian systems; any torus could be modelled. Also, Fourier series are not the only way of describing the points on a torus: one could also use other suitable cyclic tailored functions such as splines or wavelets.

Acknowledgements. This work was supported by the Academy of Finland (the project "Inverse problems of regular and stochastic surfaces" and the centre of excellence in inverse problems).

Appendix A. **Action-grid probing algorithm.** Let $J_{(m)}$, $m \in \mathbb{N}^n$ be a point in the action grid in such a way that $J_h = m_h \Delta J_h$, $h = 1, \ldots, n$, and ΔJ_h is the interval between grid points. We denote the whole action grid as a set $J_{(X)}$, $X \subset \mathbb{N}^n$. Suppose that we have a function `nearestPoint` which returns the closest grid point in $J_{(X)}$ to a given set of actions. We will also need `adjacentPoints` which returns the adjacent and diagonally adjacent points in $J_{(X)}$ to a given grid point.

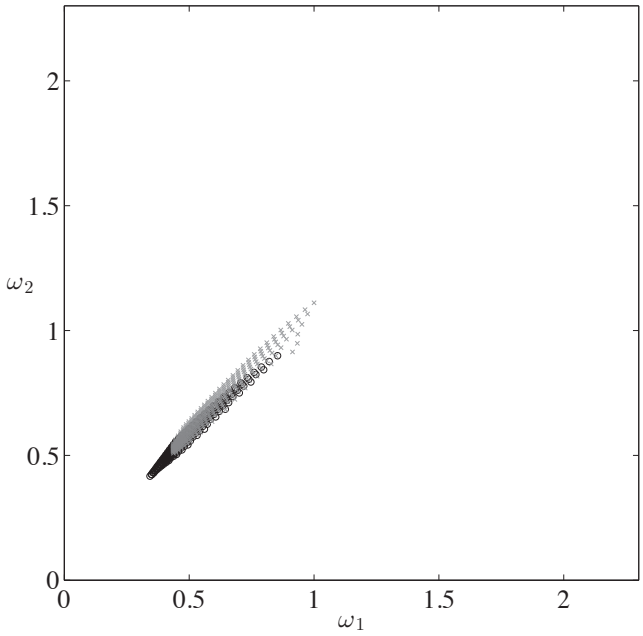

FIGURE 12. The constructed logarithmic frequencies, corresponding to the orbits in Fig. 11.

During the probing algorithm, we remove elements from $J_{(X)}$ and put them into new sets $J_{(S_0)}, J_{(S_1)}, J_{(S_2)}, \ldots$ which contain the action labels of successfully constructed tori. Let $T_{(m)}$, $m \in X$ be a torus (a set of Fourier coefficients) obtained by running the torus construction with $J_{(m)}$ given as a label. Suppose that this is implemented as a function `constructTorus` which also takes a second argument, an initial torus.

A pseudocode for the probing algorithm is given in figure 13. As input, the algorithm requires a torus T_* and a corresponding action label J_*. These are given by the unlabelled torus construction, for each potential-orbit-family pair.

REFERENCES

[1] James Binney and Sanjiv Kumar, *Angle variables for numerically fitted orbital tori*, Mon. Not. R. Astron. Soc. **261** (1993), 584–592.
[2] James Binney and David Spergel, *Spectral stellar dynamics*, Astrophys. J. **252** (1982), 308–321.
[3] ———, *Spectral stellar dynamics - II. The action integrals*, Mon. Not. R. Astron. Soc. **206** (1984), 159–177.
[4] James Binney and Scott Tremaine, *Galactic dynamics*, Princeton University Press, 1987.
[5] Tim de Zeeuw, *Elliptical galaxies with separable potentials*, Mon. Not. R. Astron. Soc. **216** (1985), 273–334.
[6] M. Hénon, *L'amas isochrone. I.*, Ann. d'Astrophys. **22** (1959), 126–139.
[7] Mikko Kaasalainen, *Hamiltonian perturbation theory for numerically constructed phase-space tori*, Mon. Not. R. Astron. Soc. **268** (1994), 1041–1050.
[8] ———, *Construction of invariant tori in chaotic regions*, Phys. Rev. E **52** (1995), 1193–1196.

input : J_*, T_*, and $J_{(X)}$
output: $T_{(S_0)}, T_{(S_1)}, T_{(S_2)}, \ldots$

$J_{(k)} \leftarrow$ `nearestPoint(`J_*`)`;
$T_{(k)} \leftarrow$ `constructTorus(`$J_{(k)}, T_*$`)`;
if $T_{(k)}$ *is good* **then** $S_0 \leftarrow \{k\}$ **else** $S_0 \leftarrow \emptyset$;
$X \leftarrow X \backslash \{k\}$;
$i \leftarrow 0$;
while $S_i \neq \emptyset$ **do**
  $S_{i+1} \leftarrow \emptyset$;
  foreach $m \in S_i$ **do**
    $J_{(A)} \leftarrow$ `adjacentPoints(`$J_{(m)}$`)`;
    foreach $k \in A$ **do**
      $T_{(k)} \leftarrow$ `constructTorus(`$J_{(k)}, T_{(m)}$`)`;
      if $T_{(k)}$ *is good* **then** $S_{i+1} \leftarrow S_{i+1} \cup \{k\}$;
      $X \leftarrow X \backslash \{k\}$;
    end
  end
  $i \leftarrow i + 1$;
end

FIGURE 13. Action-grid probing algorithm.

[9] ———, *Dynamical tomography of gravitationally bound systems*, Inverse Probl. Imag. **2** (2008), 527–546.
[10] Mikko Kaasalainen and James Binney, *Torus construction in potentials supporting different orbit families*, Mon. Not. R. Astron. Soc. **268** (1994), 1033–1040.
[11] Teemu Laakso, *Embracing integrability in stellar and planetary dynamics*, Ph.D. thesis, University of Helsinki, 2011, http://urn.fi/URN:ISBN:978-952-10-7161-4.
[12] Teemu Laakso and Mikko Kaasalainen, *Canonical methods of constructing invariant tori by phase-space sampling*, Physica D **243** (2013), 14–19.
[13] Craig C. Martens and Gregory S. Ezra, *Semiclassical mechanics of strongly resonant systems: a Fourier transform approach*, J. Chem. Phys. **86** (1987), 279–307.
[14] Colin McGill and James Binney, *Torus construction in general gravitational potentials*, Mon. Not. R. Astron. Soc. **240** (1990), 634–645.
[15] Jordi Miralda-Escude and M. Schwarzschild, *On the orbit structure of the logarithmic potential*, Astrophys. J. **339** (1989), 752–762.
[16] Y. Papaphilippou and J. Laskar, *Frequency map analysis and global dynamics in a galactic potential with two degrees of freedom*, Astron. Astrophys. **307** (1996), 427–449.
[17] Daniel P. Petersen and David Middleton, *Sampling and reconstruction of wave-numer-limited functions in N-dimensional Euclidean spaces*, Inform. Control **5** (1962), 279–323.
[18] Henri Poincaré, *Les méthodes nouvelles de la mécanique céleste*, vol. 1, Gauthier-Villars, Paris, 1892.
[19] Stephen J. Ratcliff, Kar Man Chang, and M. Schwarzschild, *Stellar orbits in angle variables*, Astrophys. J. **279** (1984), 610–620.
[20] Douglas O. Richstone, *Scale-free, axisymmetric galaxy models with little angular momentum*, Astrophys. J. **238** (1980), 103–109.
[21] Monica Valluri and David Merritt, *Torus construction*, Galaxy Dynamics (D. Merritt, J. A. Sellwood, and M. Valluri, eds.), ASP Conference Series, vol. 182, 1999.
[22] Robert L. Warnock, *Close approximations to invariant tori in nonlinear mechanics*, Phys. Rev. Lett. **66** (1991), 1803–1806.

Received xxxx 20xx; revised xxxx 20xx.